LA-UR-04-6444

Approved for public release; distribution is unlimited.

Title: PRACTICAL APPLICATION OF UNCERTAINTY-BASED VALIDATION ASSESSMENT

Author(s): Mark C. Anderson
Jeffrey E. Hylok
Ryan D. Maupin
Amanda C. Rutherford

Submitted to: 46th AIAA/ASME/ACSE/AHS/ASC
Structures, Structural Dynamics & Materials Conference
18-21 April 2005
Austin, TX

PRACTICAL APPLICATION OF UNCERTAINTY-BASED VALIDATION ASSESSMENT

Mark C. Anderson[1], Jeff Hylok[2], Ryan Maupin[2], and Amanda Rutherford[2]

Validation of simulation results by comparison with experimental data is certainly not a new idea. However, as the capability to simulate complex physical phenomena has increased over the last few decades, the need for a systematic approach to validation assessment has become evident. Organizations such as the American Society of Mechanical Engineers (ASME) and the National Laboratories are in the process of formulating validation requirements and approaches. A typical depiction of the validation process is given in Figure 1, derived from current ASME efforts regarding computational solid mechanics.

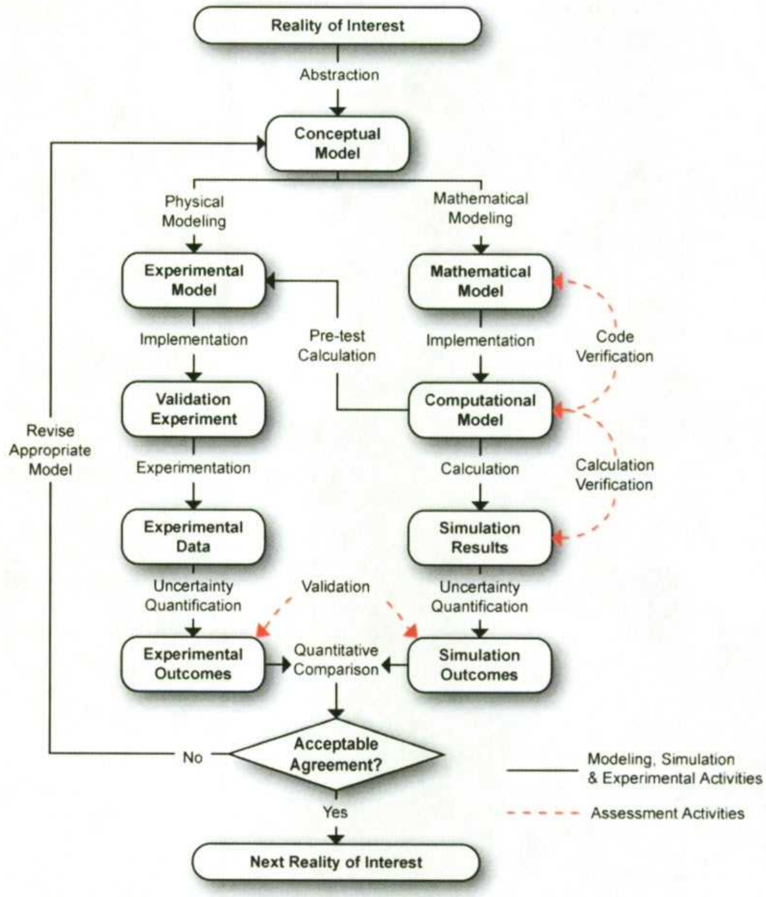

Figure 1. The verification and validation process.

---

[1] Project Leader, Weapon Response Group, Engineering Sciences and Applications Division, Los Alamos National Laboratory
[2] Technical Staff Member, Weapon Response Group, Engineering Sciences and Applications Division, Los Alamos National Laboratory

Note that uncertainty quantification plays an integral role in the validation comparison step of the process defined in the figure. This is a natural consequence of the need for verification and validation to facilitate decision-making by the customers of simulation results. Since very little is exactly known about real systems, questions of economy, reliability, and safety are best answered in the language of uncertainty.

The process illustrated in the figure above is very logical, but very general. Examples of concrete applications of this are still rare. Engineers at Los Alamos National Laboratory have been applying a systematic verification and validation process, much like that of Figure 1, to structural dynamic simulations for the past several years. These applications have resulted in the realizations that there are many details not mentioned in general process descriptions that can complicate a validation assessment. Such details include the following

- The need for a hierarchical approach in which the interactions between components within/between assemblies are considered in addition to the overall input/output behavior of the entire system.
- The need for system state and response data within the important elements of the hierarchy in addition to the observed characteristics at the system level
- Selection of appropriate response features for comparison between analytical and experimental data
- Selection of a comprehensive, but tenable set of parameters for uncertainty propagation
- Limitations of modeling capabilities and the finite element method for approximating high frequency dynamic behavior of real systems

This paper illustrates these issues by describing the details of the validation assessment for an example system. The system considered is referred to as the "threaded assembly." It consists of a titanium mount to which a lower mass is attached by a tape joint, an upper mass is connected via bolted joints, and a pair of aluminum shells is attached via a complex threaded joint. The system is excited impulsively by an explosive load applied over a small area of the aluminum shells.

The validation assessment of the threaded assembly is described systematically so that the reader can see the logic behind the process. The simulation model is described to provide context. The feature and parameter selection processes are discussed in detail because they determine not only a large measure of the efficacy of the process, but its cost as well. The choice of uncertainty propagation method for the simulation is covered in some detail and results are presented. Validation experiments are described and results are presented along with experimental uncertainties. Finally, simulation results are compared with experimental data, and conclusions about the validity of these results are drawn within the context of the estimated uncertainties.

Description of the Simulation Model

The simulation model is illustrated in Figure 2. As can be seen, the model contains a high degree of fidelity at the friction interfaces. The model contains 1.8 million nodes and 1.4 million elements. There are 496 separate contact pairs defined in the model.

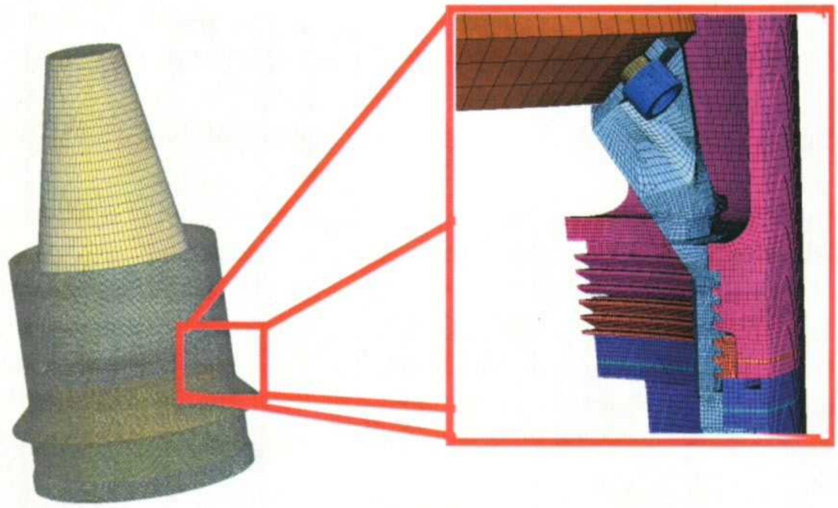

Figure 2. Simulation model showing the interfaces.

The model was run using the ParaDYN code from Lawrence Livermore National Laboratory. Preloads were applied by applying thermal loads correlated with preload test data. These loads were applied gradually to minimize vibratory response. A spatially non-uniform dynamic pressure load based on data from precision load characterization experiments was applied to excite the structure dynamically. Run time on ASCI Q is approximately 1.5 ms of dynamic simulation time per hour of clock time using 256 processors.

Feature Selection

Validating a finite element model requires choosing response features that may be derived from both experimental data and simulation results for comparison. Features should satisfy three criteria. First, they should be low-order representations of the physics of interest. In the case of the threaded assembly, the physics of interest are related to energy transmission through the structure. Second, features should have a physical meaning that can easily be explained to peers. Most importantly, features should be related to what the customer needs from the simulation.

A set of features for validation comparison of the threaded assembly simulation was chosen based on customer requirements for modeling and prediction [1]. These features are summarized in Table 1 and were evaluated at each sensor location as appropriate.

The time of arrival of signal information is helpful to identify the transmission of energy into the components in the assembly, as well as helping to confirm the elastic wave speeds of the materials. Time of arrival will be computed by determining when the dynamic acceleration, velocity, or strain signal first exceeds a prescribed threshold value. The peak value provides an indication of how much of the shock is being imparted to the system. The time at which the peak value occurs is also considered a feature of interest, similar to the time of arrival for the signal.

Temporal moments are scalar-valued features of the time history that are roughly analogous to the statistical moments of a signal. The temporal moments, $M_i(t_s)$ are calculated as weighted summations of the time signals, $y(t)$, squared, via

$$M_i(t_s) = \int_{-\infty}^{+\infty} (t - t_s)^i (y(t))^2 dt \tag{1}$$

where t_s denotes a shift in time and the subscript "i" represents the moment's order. For simplicity, the temporal moments are denoted M_i when the time shift is set to zero ($t_s=0$). The temporal moments are described in more detail in [2].

Table 1. Response features for validation assessment.

Acceleration Features	**Velocity Features**	**Strain Features**
Time of arrival	Time of arrival	Time of arrival
Peak acceleration	Peak velocity	Peak strain
Time of peak	Time of peak	Time of peak
0^{th} order temporal moment	0^{th} order temporal moment	0^{th} order temporal moment
1^{st} order temporal moment	1^{st} order temporal moment	1^{st} order temporal moment
Shock Response Spectrum		

From the temporal moments, features can be computed that have physically intuitive meanings about the signal. The first such quantity is E, the accumulated energy of the signal, which is exactly the zeroth temporal moment, M_0. The second quantity is τ, the central time, or time at which the signal energy is equal before and after, roughly analogous to the statistical mean. This value is the ratio of the first and zeroth temporal moments: $\tau = M_1/M_0$. For this analysis only the zeroth and first temporal moments were used. Also, only the first 3 ms of the response data will be used for calculating the temporal moments.

The shock response spectrum (SRS) contains information about the frequency content of the signal as well as amplitude information. The SRS synthesizes the response of a single degree of freedom (SDOF) system to a transient shock event. It calculates the vibration environment "seen" by the SDOF component due to the transient event. The SRS is commonly used to specify transient dynamic inputs for components mounted at specific locations on an assembly.

To calculate the SRS, the transient waveform is applied to a moving base, either as a force, $f(t)$, or as an acceleration signal. The response, $x(t)$, is obtained by integrating the equation of motion,

$$m\ddot{x}(t) + c\dot{x}(t) + kx(t) = f(t) \tag{2}$$

for a specific combination of mass, m, stiffness, k, and damping, c, values. Then, the solution, $x(t)$, of (2) is reduced to a single number, for example, the root mean square (RMS) displacement or the peak acceleration. The procedure is repeated for different values of the triplet $\{m, k, c\}$ covering the frequency range of interest.

Two frequency ranges were chosen for the present study. The range from 0 to 10 kHz was chosen because it is commonly used for specification of environmental requirements of components subjected to transient dynamic inputs from explosive sources. The range from 0 to

50 kHz was chosen since previous analysis and test of the threaded assembly indicated that a significant portion of the response energy occurred above 10 kHz and below 50 kHz.

Parameter Selection

Since large finite element models typically include a large collection of parameters, a necessary first step in uncertainty quantification is to determine which of those parameters most influence the response, along with their respective distributions. Then, the uncertainty of the important parameters is propagated through a finite element or a surrogate model to estimate the uncertainty associated with the response of the system.

The threaded assembly is a complex system through which we wish to characterize energy transmission from an explosive load to an upper and lower mass. Major physical phenomena have been captured in a second-generation finite element model, capitalizing on knowledge gained in the first set of threaded assembly experiments. While major details of the model have been described in [3], it is pertinent here to note some of the major changes that were made based on what was learned from the first validation exercise conducted on the threaded assembly.

a. More friction contact interfaces, with spatial variability allowed; Data from subtests utilized.
b. Preload characterization and implementation; Data from phenomena level tests utilized.
c. Input load curve characterization; Data from loads characterization tests utilized.

The choices of subtests and model improvements made were driven by the results of the first sensitivity analysis conducted on the threaded assembly [3] and further description of these subtests can be found in [4]. These results showed that friction, preload and input load levels were important to energy transmission features. The new capabilities of the revised threaded assembly finite element model led us to revising the PIRT through the use of an energy flow diagram Figure 3. The results from the subtests were also used for determination of input parameter distributions.

Engineering Judgment and The PIRT

The initial screening of parameters consisted of identification of phenomena judged to be important to energy transmission through the threaded assembly in the form of a phenomena importance and ranking table (PIRT). The PIRT was derived from the energy flow diagram shown in Figure 2. The energy flow diagram was utilized to better understand the path of energy through the different parts of the threaded assembly. This aided in defining important interfaces, material properties, etc. Once the all-inclusive list of phenomena and properties were listed in the PIRT, engineering judgment and PIRT logic were employed to obtain a more tractable list of parameters. Other considerations were the current ability to model the phenomena described. The original table included some 50 parameters relating to friction, preload, material properties and load input to the structure. A smaller parameter set was down-selected using engineering judgment (rooted in experience with the threaded assembly problem) to only 11 parameters.

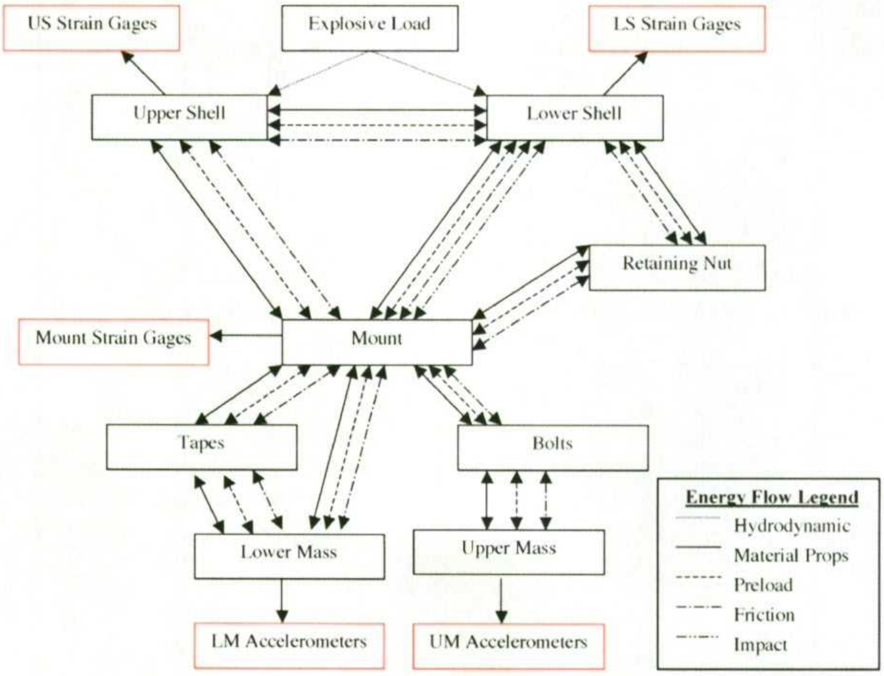

Figure 3. Energy flow diagram.

The parameters retained are shown in Table 2. We must acknowledge that anytime we reduce the input space, we accept some risk that the wrong decision was made, i.e. a parameter was identified as un-important, but actually was important, either on its own or as part of some multi-parameter interaction.

Table 2. Parameters retained from the PIRT for further study.

Parameter	Mean	St Dev	Distribution
1. Lower Shell – Upper Shell Friction	0.52	0.08	Truncated Weibull
2. Mount – Lower Shell Friction	0.47	0.05	Lognormal
3. Mount – Nut Friction Friction	0.70	0.12	Lognormal
4. Mount – Lower Mass Friction	0.44	0.10	Truncated Normal
5. Mount – Tape Friction	0.44	0.10	Truncated Normal
6. Mount – Upper Shell Friction	0.81	0.16	Truncated Normal
7. Nut – Lower Shell Friction	0.47	0.05	Lognormal
8. Lower Mass – Tape Friction	1.00	0.12	Lognormal
9. Nut Preload (model CTE)	1.1e-3	2.0e-4	Truncated Normal
10. Upper Shell Preload (model CTE)	1.5e-3	4.0e-4	Truncated Normal
11. Tape Preload (model CTE)	1.0e-2	2.0e-4	Truncated Normal

Note that the parameters retained for further study relate to the friction coefficients at various interfaces, and the preloads. Parameters that were eliminated were the geometry parameters, the material properties, and the bolt-mount interface properties and preload. Elimination of the geometry parameters and material properties was relatively low risk (in a main effects sense) since there was high confidence in their means and/or they were judged to be of low importance. Elimination of the bolt-mount interface properties and preload may involve somewhat higher risk, because the interface was not characterized fully, and it has a large effect on energy transmission to the accelerometers on the upper mass. However, in the current model form, this interface has been modeled as tied contacts and we had only one number (the specification) for the preload. This parameter may have to be revisited in the future, if the model is improved to include bolt preload.

Since they were characterized with independent tests, the input loads are not listed in the parameter table and will not be treated as parameters. Rather, sensitivity studies will be conducted at each load level to see if the sensitivities of the other parameters change with increasing load level.

Linear Sensitivity Analysis

Before proceeding to increased levels of complexity in sensitivity analysis that would require a significant number of runs of the finite element model, it was deemed prudent to consider the computational run budget available to us for all parameter screening and uncertainty quantification activities. It has been estimated that the maximum number of computational runs that could be accomplished in the two months allotted for parameter screening and uncertainty quantification activities is 200. With 11 parameters, a standard linear sensitivity analysis would consist of each parameter run at a high and a low value, with the others at their nominal values, and one run with all parameters at their nominal value, i.e., 23 runs. These runs would have to be done for each of three load levels, meaning that we would have 69 runs. This leaves 131 runs for uncertainty quantification activities.

Another option explored was a Plackett-Burman (PB) design, which is a two level design that requires just 12 runs for 11 parameters [6, 7]. With our three load cases, this would mean a total of 36 runs, leaving 164 for UQ activities. The disadvantage of using a PB design is its complex alias structure. In this design, main effects are aliased with many two-factor interactions. We conducted this kind of study recognizing that the information that we obtained might be aliased with higher order effects.

Based on the criterion of retaining any parameter that was significant for one or more features, no parameters could be eliminated from consideration after the linear sensitivity analysis. Thus, all 11 parameters were retained for the uncertainty analysis.

Simulation Uncertainty Analysis

With 164 runs remaining after the PB sensitivity analysis, a couple of options were available for the propagation of uncertainties. The first option considered was the Latin Hypercube Sampling (LHS) scheme. This involves discretizing each of the distribution functions for the 11 input parameters, randomly pairing the samples from each of the bins, and calculating the output features for these pairings. With 164 runs remaining, we could perform 50 LHS runs (meaning the input parameter distributions would be divided into 50 equal probability intervals) at each of the three remaining load levels.

One advantage of using the LHS scheme is that the need for a surrogate model is avoided, since the finite element model is sampled directly. Another advantage is that the number of samples chosen depends only on how finely binned the distributions are, not on how many input parameters there are. However, while it is likely that 50 samples would be enough, a preliminary search of the literature has not yielded a method for *a priori* calculation of convergence of the output feature cumulative distribution function (CDF) derived. Helton and Davis [5] suggest a method for determining confidence intervals on the output feature CDF, but it is calculated after the LHS runs have been completed. Also, sensitivity analysis with LHS runs must be conducted through fitting models (though it is unclear what the alias structure would be, due to inherent lack of orthogonality), use of scattergrams, etc. Because of these difficulties and because of previous success with other methods, a Latin Hypercube sampling scheme was not used.

Instead, a response surface analysis method was chosen for uncertainty quantification. To form a response surface with 11 input parameters, we will use a central composite design (CCD) to determine the needed simulation runs [6,7]. The CCD may be thought of as a 2-level fractional factorial design, appended with runs at the axial points (these look like the low-nominal-high linear sensitivity analysis). With about 55 runs available for each load level model, we considered a $1/64^{th}$ fraction design for the 2-level portion (32 runs), plus the axial points (22 runs), plus the center point. Care was taken when choosing which parameters the model should retain, so that the alias structure is capable of capturing the interactions between frictions and preloads that were anticipated to be important. An advantage of using the response surface analysis is that an analysis of variance will fall out of the model chosen and we will be able to conduct a Level 3 sensitivity analysis as a byproduct of the design.

Note that any reduction in the number of parameters that will be taken at the Level 2 stage increase the resolution of the 2-level fractional factorial part of the response surface analysis, meaning that we will be able to capture additional higher order effects with less aliasing. Once the response surface is fit for each load level, it will be Monte Carlo sampled (because evaluation time for the polynomial is negligible), and a distribution on the output features will be formulated.

Evaluation of the uncertainty on the simulation output due to parametric uncertainty is still in progress. If the schedule permits, a comparison of the response surface approach and one based on LHS will be conducted for at least one load level.

Description of the Experiments

To provide data for validation comparisons and predictive accuracy estimation, twelve assembly-level tests are planned. Previous assembly-level testing indicated that control of parameters significant for the numerical model, other than the load level, was not precise enough to provide useful comparison data. These earlier experiments and the subsequent analysis showed that replication of experiments to facilitate estimation of experimental uncertainties, along with a rigorous characterization of the as-built test hardware, would be required to permit meaningful comparisons with model predictions.

The test matrix for the current series of tests is shown in Table 3. Three load levels, designated *low*, *medium*, and *high*, were planned based on charge development efforts. In addition to the controlled variation of load level, four sets of upper and lower shells were used for the tests. The separate sets of shells permitted variation of system geometry in an uncontrolled, but measurable manner. Thus, there were four tests at each load level with each of the four replicates using a different set of shells as shown in the table. Note that there was also an uncontrolled variation of friction characteristics in the threaded joint. This variation, along with that of preload, was handled using the results of friction and preload testing on the independent of the validation experiments.

Table 3. Validation test matrix.

	Load Level		
	Low	Medium	High
Shell Set	1	1	1
	2	2	2
	3	3	3
	4	4	4

A typical test setup is depicted in Figure 4. The test article was suspended as a pendulum from a single point to permit overall impulse response to be estimated by fiber optic "light ladder" data. The explosive load was comprised of strips of Primasheet over a neoprene patch. Simultaneous detonation of the explosive strips was accomplished via the use of an explosive timing manifold, or lens, of the same Primasheet material. Tri-axial accelerometers were mounted on the upper and lower surfaces of the upper and lower masses. The shells and mount were instrumented with 38 strain gages. All data were recorded with Nicolet oscilloscopes and Odyssey data acquisition systems. Data were recorded at 10 or 20 Mhz, depending on the channel and oscilloscope.

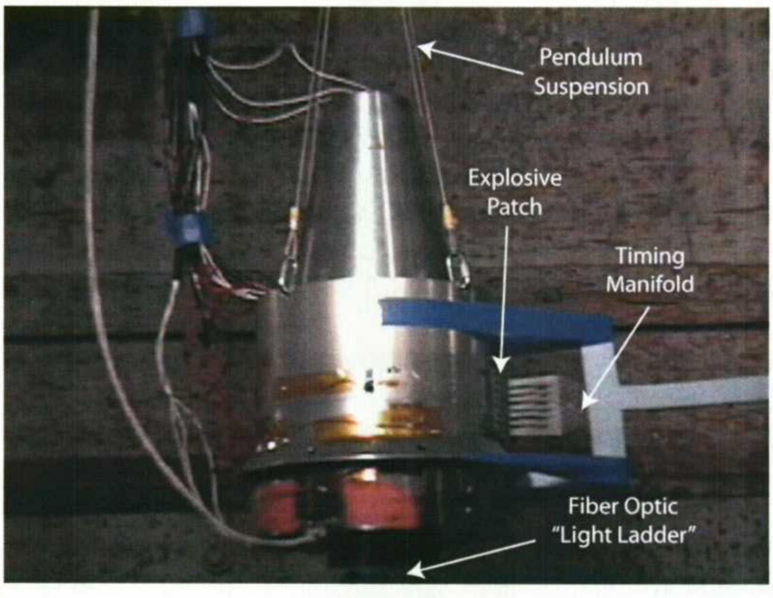

Figure 4. Validation experiment test setup.

Data processing is currently under way. Results will be available for the full paper.

Experimental Uncertainty Analysis

The uncertainty analysis of the experimental results was based simply on calculating the first and second order statistics (means and standard deviations) of the relevant response features for each sensor, and for both frequency ranges of interest.

To facilitate comparison with simulation results, the experimental data and simulation results were processed to put them on a common basis. The measured response time histories were temporally synchronized with the simulation time histories by determining a common "zero" time and shifting the data accordingly. Since the data were sampled at a higher rate (10 to 20 Mhz) than the simulation results (1 Mhz), the data were re-sampled to the analysis rate after Nyquist filtering. Filtering to the frequency ranges of interest (0-10 kHz and 0-50 kHz) was accomplished using identical algorithms for the data and simulation results.

The uncertainty analysis of the data is in progress. Final results will be available for the full paper.

Preliminary Results

Data processing and uncertainty analysis is still under way. Some test results have been extracted from the high load level test on the third shell set for comparison with baseline (nominal) calculations. These results are shown in Figure 5. Depicted in the figure are the peak strain and 0^{th} order temporal moment of strain for selected strain gages. Responses for this comparison were filtered at 50 kHz. The test data include the mean and one standard deviation points for each measurement. Note that the analytically predicted peak strains are significantly different from the measured strains. This result is to be expected for high frequency "point" data since most finite element simulations are still incapable of accurately predicting point quantities at higher frequencies. In contrast, the comparison for the 0^{th} order temporal moment is much better. The latter feature is integrated in the temporal sense, obviating the details of the frequency content in both the data and the simulation results. It is anticipated that, for the foreseeable future, spatially and/or temporally integrated features will be more useful for comparing high frequency response calculations with data.

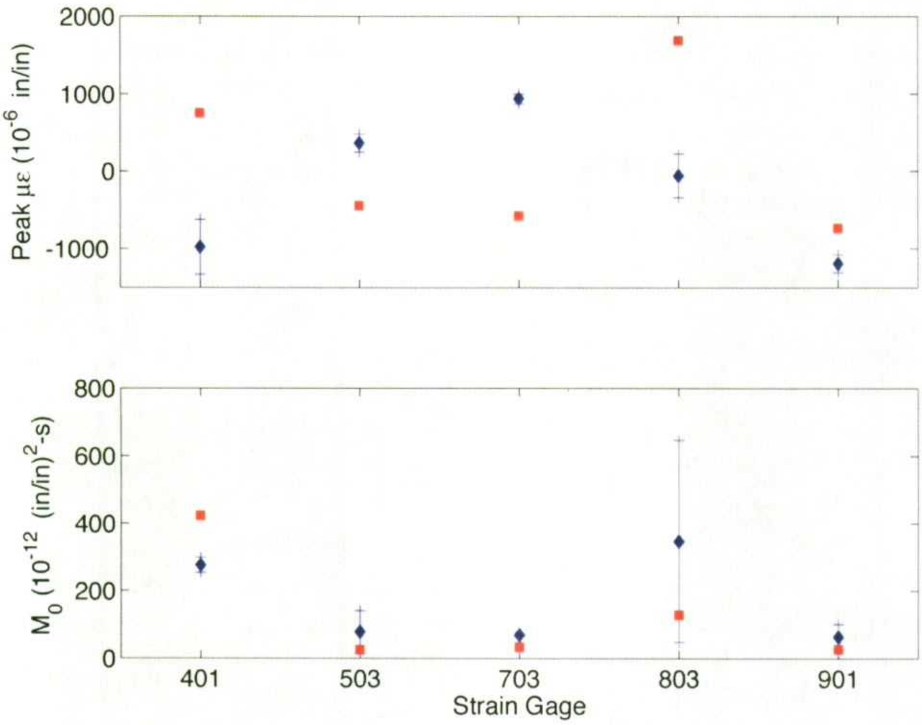

Figure 5. Comparison of baseline simulations with experimental data.

Complete results will be available for the full paper.

References

1. Stevens, B., "Customer Requirements for the ASC Threaded Assembly V&V Milestone," *Los Alamos National Laboratory Memorandum ESA-WR: 04-080.*
2. Smallwood, D. O., "Characterization and Simulation of Transient Vibrations Using Band Limited Moments," *Shock and Vibration,* Vol. 1, No. 6, 1994, pp. 507–527.
3. Doebling, S.W., et al. "Validation of the Transient Structural Response of a Threaded Assembly: Phase 1," *Los Alamos National Laboratory Publication LA-14104-MS*, 2004.
4. Maupin, R.D. "Threaded Assembly Pre-Shot Calculation Report," *Los Alamos National Laboratory Publication LA-UR 04-1874*, 2004.
5. Helton, J.C. and Davis, F.J. "Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems," *Reliability Engineering and System Safety*, Vol. 81, 2003, pp 23-69.
6. Meyers, R.H. and Montgomery, D.C., *Response Surface Methodology*, Wiley Series in Probability and Statistics, New York, 1995.
7. Stat Ease, *Design Expert 6.0.3* Software, Minneapolis, 2000